

COMPLEX SIGNAL DECOMPOSITION AND MODELING

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of priority under 35 U.S.C. § 119(e) to U.S.

5 Provisional application serial no. 60/187,949 filed March 9, 2000.

BACKGROUND OF THE INVENTION

1. Field of the Invention

10 The present invention relates to process modeling and monitoring, especially using complex or periodic signals. Moreover, the invention relates to complex signal decomposition, modeling and classification for use in monitoring the operational state of any machine, process or signal. More particularly, a complex signal can be classified using the present invention for health monitoring or for information rectification.

15

2. Description of the Related Art

Well known conventional modeling and analysis methods use several sensors measuring operational system parameters to monitor equipment or processes - generically "systems". The related signals from sensors can be observed directly to understand how the system is functioning. Alternatively, for unattended operation, methods are well known for comparing real-time sensor signals against stored or predetermined thresholds in an automated fashion. When the signals exceed those thresholds, an exception condition or an alarm is generated, thus requiring human intervention only when a sensor datum value exceeds a corresponding threshold.

20 Such methods focus on using the instantaneous value of sensors and other parameters to describe the current state of a system, but do not take advantage of time-domain information locked up in the sensor signals. It would be useful to take advantage of such time-domain information to better monitor the system, and even determine what operational state (among many acceptable states) the system is in.

25

In the field of vibration analysis, methods are known for examining the power spectral density function from an accelerometer or acoustic pickup to provide means for monitoring rotating or cyclic equipment. Typically, frequencies of interest are examined, and thresholds (lower or upper limit) are placed on the power level expected for these frequencies. If a threshold is pierced, this is indicative of an unsatisfactory operating condition or a developing problem. A great deal of work is involved in identifying the frequencies of interest and expected power levels for each particular piece of equipment that is monitored in this fashion. Problem diagnosis is also typically very specific to the kinds of indications presented with the appearance of the particular problem, and must be worked out specifically for each machine. It would be useful to have an empirical data-driven way of determining the health or the operational state of a machine based on one or more vibration or acoustic signals.

In a different technical area, digital data transmission is frequently accomplished – whether over a cable (e.g. Cat. 5, coaxial cable, etc.) or through radio transmission (e.g. broadcast, digital telecommunication, an IEEE 802.11b interface) – by modulation of an analog carrier signal. Further, to improve data transmission rates, the data being transmitted is compressed and encoded onto the transmission signal carrier, typically as sinusoidal waves encoding binary data in the phase and amplitude of the wave. Presently, well-known data encoding and transmission techniques include quadrature amplitude modulation (QAM) and discrete multitone (DMT). Well-known methods for extracting such encoded data include frequency filtering, signal decomposition and wavelet analysis.

However, during transmission these types of signals can suffer from attenuation and interference due to noise or transmission media deterioration, for example. In some cases, noise and signal degradation is sufficient to all but obliterate the original transmitted signal, making it impossible to extract the data encoded therein using prior art techniques. Accordingly, when noise or degradation is high, it would be useful to be able to reconstruct a meaningful signal from the

noisy and/or attenuated signal that is being received. This essentially amounts to determining which of a finite set of datagrams an attenuated signal most closely resembles. There is a need for a signal analysis method that may be applied to a single complex signal to extract an original signal.

5 One empirical model-based monitoring technique known in the art is described in U.S. Patent No. 5,764,509 to Gross et al., the teachings of which are incorporated herein by reference. In this technique, multiple sensor signals measuring physically correlated parameters are modeled in an empirical technique to provide estimates of those values. Discrepancies between the estimates and the
10 actual values from the sensors indicate a developing process or machine failure, or sensor failure. The model generates the estimates using a reference library of selected historic snapshots of sensor values representative of known operational states. However, the described embodiments therein do not utilize the time domain information in the sensor signals, and instead usually treat the data in distinct and
15 disconnected time-contemporaneous snapshots. It would be useful to provide the kind of empirical modeling of Gross et al. for use with time domain information. What is needed is a way of using a complex signal as an input to a multivariate modeling system such as that of Gross et al.

Where time domain information is locked up in one or more sensor or
20 parameter signals detected from an instrumented process or machine, what is needed is a way to model the process or machine with the time-domain signal for one or more acceptable and identifiable states of operation, and to do so without investing a great deal of time and effort in coming up with first-principles equations that approximate those states. What is further needed is a way to categorize or
25 classify system operational states based on a complex signal.

SUMMARY OF THE INVENTION

The present invention achieves the above needs by using an empirical modeling engine in combination with a complex signal decomposition technique,

extracting multivariate information from at least one single complex signal. The extracted components are provided as individual inputs to the empirical modeling engine. The empirical modeling engine compares the extracted component inputs against expected values to derive more information about the actual or originating
5 signal or about the state of the system generating the signal.

A complex signal or parameter, such as a time varying electrical signal (e.g., a electrocardiogram), is decomposed into multiple related signals by band notch filtering, wavelet analysis or other means. The multiple signals essentially form multivariate inputs to an empirical modeling engine, which generates signal
10 estimates in response to the inputs. The modeling engine utilizes a stored reference set of prior known multivariate snapshots to generate the estimates. The estimates and the actual signals are compared in a sensitive statistical test to detect deviations between them, indicative of an unrecognized operational state. Alternatively, the empirical modeling engine compares the multivariate input snapshot to the
15 snapshots in the reference set to determine which reference set snapshots are most similar to the input. A classification associated with the stored reference snapshots having the highest similarity is selected as the classification of the inputs, and the operational state represented by the complex signal is thus determined from a finite set of such known states represented in the reference set. Accordingly, not only can
20 operational states of a system providing the complex signal be categorized, but also original datagrams can be reconstructed from degraded encoding carrier waves for digital data transmission, thereby affording greater range and accuracy.

Briefly summarized, the invention is implemented in a computer or other processor, having a memory for storing the reference set of snapshots of related
25 signal values, and associated classifications or states, determined from empirical prior observations. An input means can be a data bus, messaging network or direct data acquisition device. The processor is disposed to compute estimates and compare them with the actual inputs to determine differences, or is alternatively disposed to select the classification from the reference set with highest similarity to

the inputs. Classification or difference detection can be output to a screen or data file; used to send a paging message, email or fax; or made available to downstream data processing applications in the same computer or another computing system sharing a network or bus connection.

5

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as the preferred mode of use, further objectives and advantages thereof, is best understood by reference to the following detailed description of the embodiments in conjunction with the accompanying drawings, wherein:

10

FIG. 1 is a diagram of the general embodiment of the present invention for signal decomposition and empirical modeling;

15

FIG. 2 is a diagram of an embodiment of the present invention for operational deviation detection;

FIG. 3 is a diagram of an embodiment of the present invention for decoding digital information that has been compressed and encoded onto a modulated analog carrier wave that has been attenuated or degraded;

20

FIG. 4 illustrates 16 predetermined composite waveforms for conveying encoded digital information;

FIG. 5 is an example of a discrete wavelet transform on the Signal Number 4 shown in FIG. 4 using the Haar wavelet;

FIG. 6 shows the 16 composite waveforms of FIG. 4 with random additive noise to emulate distortion in transmission;

25

FIG. 7 shows bar graphs indicating similarity values of the 16 signatures of FIG. 4 as compared with the noisy signals of FIG. 6; and

FIG. 8 is a flow diagram of the preferred embodiment signal comparison method.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Turning now to the drawings and, more particularly, FIG. 1 generally shows the preferred embodiment signal decomposition and modeling system 100 of the present invention, which generally includes a signal receiver 108, a signal decomposition module 102, a reference library 104 and an empirical modeling engine 106. Signal receiver 108 serves to provide the complex signal to the signal decomposition module 102, which converts the signal to a plurality of correlated component signals. Empirical modeling engine 106 is disposed to receive the plurality of inputs, and with reference to a set of data embodying classification information or known operational information stored in reference library 104, output an indication about the complex signal and the system it is drawn from.

System 100 has application to any complex signal, whether naturally occurring, derived from a system, system operation or a process being carried out in a system or otherwise generated, naturally or artificially and received at the signal decomposition and modeling system 100 by an appropriate receiver 108. In particular, the complex signal may be representative of a physical parameter of a system, process or machine, including, for example, a biological process or system. More particularly, the complex signal can be an electrocardiogram signal for monitoring a heart; the current signature on the power draw of an electric motor; the pressure transducer signal from a metal stamping machine; and so on. Further, the complex signal may be a transmission carrier signal with information encoded thereon such as from telecommunications transmission, e.g., quadrature amplitude modulation (QAM) or discrete multitone (DMT) transmission. The above-described examples of complex signals are provided for example only and not intended as a limitation.

signal. Alternatively, the complex signal may be decomposed using frequency filters to extract an individual signal component within the bandpass of each frequency filter. The signal may be digitized before or after decomposition, with signal components being a digital result of decomposition. Thus, the output of the
5 signal decomposition module 102, Y_{input} , is multiple correlated signal components. Each of these correlated components from ones of the frequency filters, for example, may be represented as a coefficient for a corresponding frequency, the coefficients for all of the frequencies forming the input vector, Y_{input} .

The reference library 104 generally contains prior data that is characteristic of
10 expected signal component values. This prior data can be distilled from data collected off a real process or machine operating in known states. Alternatively, in the case of rectifying degraded encoded communications signals, the prior data in the reference library can be exemplary datagrams generated by algorithm. In either case, the reference library comprises at least a plurality of snapshots or vectors, each
15 vector containing a like number of elements. Each element is the value of a component signal comprising the complex signal. Thus, if a complex signal is decomposed in module 102 using frequency band filters, then a given vector could comprise the amplitudes of each of the component frequency band signals at a single observation of those component signals, and would have as many elements as
20 there are component signals. If the complex signal is decomposed using discrete wavelet transform, the reference set vectors can be constituted a number of alternative ways; for example by using the coefficients of a selected level of discrete wavelet transform as the vector element values. In any case, according to the invention, the reference library preferably comprises vectors having at least two
25 elements. Furthermore, all the vectors must have a same number of elements. Finally, the elements in each vector must come from the same corresponding source, e.g., the first element of each vector comes from a first component signal, and the second element is always the value of the second component signal, and so on.

Typically, historical signal component values (signature vectors) are maintained in matrix form (designated **D** herein) in the reference library 104. Each signature vector is a column (or row) of the history matrix **D** and represents a previous decomposition of the complex signal at a selected period, i.e., is a snapshot of the signal and contains an entry for each component comprising the complex signal (see below for further discussion). The set of elements in the matrix **D** are referred to herein as a signal signature set. Snapshots (signature vectors) in the signal signature set of reference library 104, typically, are derived from nominal signals and so, are representative of expected complex signal states.

The empirical modeling module 106 periodically receives the decomposed signal samples (i.e., signal components) \mathbf{Y}_{input} as signal snapshots or data frames. Each snapshot \mathbf{Y}_{input} is compared to the data snapshots in the matrix **D** in the reference library 104. The empirical modeling module 106 compares the current signal snapshot \mathbf{Y}_{input} received from the decomposition module 102 with reference set snapshots for "similarity". This measure of "similarity" is computed using a similarity operator. According to the invention, the similarity operation for a pair of snapshots or vectors being compared returns a value, typically between zero and one, where zero represents dissimilarity and one represents completely identical snapshots. According to one class of similarity operator that works on an element-by-element basis, the measure of similarity is computed according to:

$$\vec{S} = \vec{D}^T \otimes \vec{Y}_{input} \quad (1)$$

where **D** is the reference library set of snapshots, \mathbf{Y}_{input} is the input snapshot, **S** is the vector of all snapshot similarities (one for each comparison of the \mathbf{Y}_{input} to a reference snapshot) and the similarity operation is designated by the symbol \otimes .

The similarity operation can be selected from a variety of known operators that produce a measure of the similarity or numerical closeness of rows of the first operand to columns of the second operand. The result of the operation as generally applied to two matrix operands is a similarity matrix wherein the similarity value of

the *i*th row and *j*th column is determined from the *i*th row of the first operand and the *j*th column of the second operand (the operand above being a vector Y_{input} and having only one column). The resulting element (*i,j*) is a measure of the sameness of these two vectors. In the present invention, the *i*th row of the first operand generally
5 has elements corresponding to temporally related component values for a given decomposition of the complex signal from the process or machine, and the same is true for the *j*th column of the second operand. Effectively, the resulting array of similarity measurements represents the similarity of a reference vector in one operand to a current input vector in the other operand.

10 By way of example, one similarity operator that can be used compares the two vectors (the *i*th row and *j*th column) on an element-by-element basis. Only corresponding elements are compared, e.g., element (*i,m*) with element (*m,j*) but not element (*i,m*) with element (*n,j*). For each such comparison, the similarity is equal to the absolute value of the smaller of the two values divided by the larger of the two
15 values. Hence, if the values are identical, the similarity is equal to one, and if the values are grossly unequal, the similarity approaches zero. When all the elemental similarities are computed, the overall similarity of the two vectors is equal to the average of the elemental similarities. A different statistical combination of the elemental similarities can also be used in place of averaging, e.g., median.

20 Another similarity operator that can be used in the present invention is the bounded area ratio test (BART) described in U.S. Patent No. 5,987,399, which is incorporated by reference. BART is a prior art similarity operator, wherein an internal angle is used to gauge the similarity of two values. A right triangle is formed for each signal component with the base (hypotenuse) of each right triangle
25 bounded by an expected magnitude range over all snapshots in the signal signature set for the particular component. The right angle vertex is preferably located at a point above the median or mean of the range, and at a height *h* that forms the right angle, the right angle vertex being the apex of the right triangle. At each comparison during system monitoring, BART maps two points X_1 and X_0 to the base; one point

representative of an expected component value and the second point is a current component value. These two points are located on the base according to their magnitude within the range of values in the signal signature set. An internal comparison angle θ is formed at the apex above the base by drawing a line to the apex from each of the mapped points:

$$\theta = \tan^{-1}\left(\frac{h}{X_1}\right) - \tan^{-1}\left(\frac{h}{X_0}\right) \quad (2)$$

The internal angle is the basis by which two values are compared for similarity, i.e., identical points result in a 0° angle and completely dissimilar points result in a right angle. Then, the elemental similarity for the i^{th} element is:

$$s_i = 1 - \frac{\theta_i}{\pi/2} \quad (3)$$

As indicated above, the elemental similarities can be statistically averaged or otherwise statistically treated to generate an overall similarity of a snapshot to another snapshot, as is called for according to the invention.

Yet another class of similarity operator that can be used in the present invention involves describing the proximity of one signature vector to another vector in n -space, where n is the dimensionality of the vector of the current snapshot of the monitored process or machine. If the proximity is comparatively close, the similarity of the two vectors is high, whereas if the proximity is distant or large, the similarity diminishes, ultimately vanishing. By way of example, Euclidean distance between two vectors can be used to determine similarity. In a complex signal decomposed into 20 components for example, the Euclidean distance in 20-dimensional space between the currently monitored snapshot, comprising a 20-element vector, and each vector in the reference library provides a measure of similarity, as shown:

$$S = \frac{1}{1 + \frac{\|\vec{x} - \vec{d}\|^\lambda}{c}} \quad (4)$$

wherein \mathbf{X} is the current snapshot, and \mathbf{d} is a vector from the reference library, and λ and c are user-selectable constants.

It should be understood that, although specific similarity operators are described herein, any general similarity operator may be used in the empirical modeling engine of the present invention. The output of the empirical modeling engine 106 of FIG. 1 can be an indication based on the similarity measures generated for each comparison of a reference library vector to the input vector. A number of alternatives are discussed below.

Turning now to FIG. 2, an embodiment of the invention is shown for use in monitoring the operation of a process or machine, for detection of incipient failures, sensor failure, or operational state change. Items identical to those of FIG. 1 are labeled identically. The empirical modeling engine 106 generates a snapshot $\mathbf{Y}_{\text{expected}}$, of estimated or expected values in response to the input snapshot $\mathbf{Y}_{\text{input}}$, based on the similarity measures with the reference library 104 as discussed below.

The current decomposed snapshot $\mathbf{Y}_{\text{input}}$ is subtracted from the estimate $\mathbf{Y}_{\text{expected}}$ in an adder 110 to yield component residual values. When the components of the decomposed snapshot $\mathbf{Y}_{\text{input}}$ and the expected snapshot $\mathbf{Y}_{\text{expected}}$ are close, the residuals are very nearly zero, indicating no incipient failures or state changes. When they are different, the residuals are other than zero. Using a sophisticated test, this difference can be ascertained even though the residuals are within the noise band of the system. A test module 112 is provided to receive and test the residuals to determine at the earliest determinable observation whether a deviation is evidenced in the residuals or not. There is generally one residual for each component signal in the model. An indication of difference between the expected

values and the actual values of the component signals points to a deviation in the operation of the monitored process or machine as measured using the complex signal input to receiver 108. Finally, a diagnostic lookup table 114 can optionally be provided to output a likely failure mode is response to the pattern or deviations
5 detected by the test module 112 in the residuals.

The empirical modeling engine 106 determines expected values of the complex signal components according to the following equation:

$$\vec{y}_{Expected} = \overline{D} \bullet \vec{W} \quad (5)$$

where \mathbf{W} is a weight vector having as many elements N as there are columns (or snapshots) in the reference library \mathbf{D} 104, generated by:

$$\vec{W} = \frac{\hat{\vec{W}}}{\left(\sum_{j=1}^N \hat{W}(j) \right)} \quad (6)$$

$$\hat{\vec{W}} = \left(\overline{D}^T \otimes \overline{D} \right)^{-1} \bullet \left(\overline{D}^T \otimes \vec{y}_{input} \right) \quad (7)$$

10 where the similarity operation is represented by the circle with the cross-hatch inside it. The superscript "T" here represents the transpose of the matrix, and the inverse of the matrix or resulting array is represented by the superscript "-1". Importantly, there must be row correspondence to like signal components for the rows in \mathbf{D} , \mathbf{Y}_{input} and $\mathbf{Y}_{expected}$. That is, if the first row of the reference matrix \mathbf{D}
15 corresponds to values for a first component of the decomposed complex signal, the first element of \mathbf{Y}_{input} must also be the current value (if operating in real-time) of that same first component.

Test module 112 can implement a comparison of the residuals to selected thresholds to determine when an alert should be output of a deviation in the
20 complex signal from recognized states stored in the reference library. Alternatively, a statistical test, preferably the sequential probability ratio test (SPRT) can be used to

determine when a deviation has occurred. The basic approach of the SPRT technique is to analyze successive observations of a sampled parameter. A sequence of sampled differences between the generated expected value and the actual value for a monitored component signal should be distributed according to some kind of distribution function around a mean of zero. Typically, this will be a Gaussian distribution, but it may be a different distribution, as for example a binomial distribution for a parameter that takes on only two discrete values (this can be common in telecommunications and networking machines and processes). Then, with each observation, a test statistic is calculated and compared to one or more decision limits or thresholds. The SPRT test statistic generally is the likelihood ratio l_n , which is the ratio of the probability that a hypothesis H_1 is true to the probability that a hypothesis H_0 is true:

$$l_n = \frac{(y_1, y_2, \dots, y_n | H_1)}{(y_1, y_2, \dots, y_n | H_0)} \quad (8)$$

where Y_n are the individual observations and H_n are the probability distributions for those hypotheses. This general SPRT test ratio can be compared to a decision threshold to reach a decision with any observation. For example, if the outcome is greater than 0.80, then decide H_1 is the case, if less than 0.20 then decide H_0 is the case, and if in between then make no decision.

The SPRT test can be applied to various statistical measures of the respective distributions. Thus, for a Gaussian distribution, a first SPRT test can be applied to the mean and a second SPRT test can be applied to the variance. For example, there can be a positive mean test and a negative mean test for data such as residuals that should distribute around zero. The positive mean test involves the ratio of the likelihood that a sequence of values belongs to a distribution H_0 around zero, versus belonging to a distribution H_1 around a positive value, typically the one standard deviation above zero. The negative mean test is similar, except H_1 is around zero minus one standard deviation. Furthermore, the variance SPRT test can be to test

whether the sequence of values belongs to a first distribution H_0 having a known variance, or a second distribution H_2 having a variance equal to a multiple of the known variance.

For residuals derived for component signals from a complex signal behaving as expected, the mean is zero, and the variance can be determined. Then in run-time monitoring mode, for the mean SPRT test, the likelihood that H_0 is true (mean is zero and variance is σ^2) is given by:

$$L(y_1, y_2, \dots, y_n | H_0) = \frac{1}{(2\pi\sigma)^{n/2}} e^{\left[-\frac{1}{2\sigma^2} \sum_{k=1}^n y_k^2 \right]} \quad (9)$$

and similarly, for H_1 , where the mean is M (typically one standard deviation below or above zero, using the variance determined for the residuals from normal operation) and the variance is again σ^2 (variance is assumed the same):

$$L(y_1, y_2, \dots, y_n | H_1) = \frac{1}{(2\pi\sigma)^{n/2}} e^{\left[-\frac{1}{2\sigma^2} \left(\sum_{k=1}^n y_k^2 - 2 \sum_{k=1}^n y_k M + \sum_{k=1}^n M^2 \right) \right]} \quad (10)$$

The ratio l_n from equations 9 and 10 then becomes:

$$l_n = e^{\left[-\frac{1}{2\sigma^2} \sum_{k=1}^n M(M - 2y_k) \right]} \quad (11)$$

A SPRT statistic can be defined for the mean test to be the exponent in equation 11:

$$SPRT_{mean} = -\frac{1}{2\sigma^2} \sum_{k=1}^n M(M - 2y_k) \quad (12)$$

The SPRT test is advantageous because a user-selectable false alarm probability α and a missed alarm probability β can provide thresholds against which $SPRT_{mean}$ can be tested to produce a decision:

1. If $SPRT_{mean} \leq \ln(\beta / (1-\alpha))$, then accept hypothesis H_0 as true;
2. If $SPRT_{mean} \geq \ln((1-\beta) / \alpha)$, then accept hypothesis H_1 as true; and

3. If $\ln(\beta/(1-\alpha)) < \text{SPRT}_{\text{mean}} < \ln((1-\beta)/\alpha)$, then make no decision and continue sampling.

For the variance SPRT test, the problem is to decide between two hypotheses: H_2 where the residual forms a Gaussian probability density function with a mean of zero and a variance of $V\sigma^2$; and H_0 where the residual forms a Gaussian probability density function with a mean of zero and a variance of σ^2 . The likelihood that H_2 is true is given by:

$$L(y_1, y_2, \dots, y_n | H_2) = \frac{1}{(2\pi V^{1/2} \sigma)^{n/2}} e^{\left[-\frac{1}{2V\sigma^2} \sum_{k=1}^n y_k^2 \right]} \quad (13)$$

The ratio l_n is then provided for the variance SPRT test as the ratio of equation 13 over equation 9, to provide:

$$l_n = V^{-1/2} e^{\left[-\frac{1}{2\sigma^2} \sum_{k=1}^n y_k^2 \left(\frac{1-V}{V} \right) \right]} \quad (14)$$

and the SPRT statistic for the variance test is then:

$$\text{SPRT}_{\text{variance}} = \frac{1}{2\sigma^2} \left(\frac{V-1}{V} \right) \sum_{k=1}^n y_k^2 - \frac{\ln V}{2} \quad (15)$$

Thereafter, the above tests (1) through (3) can be applied as above:

1. If $\text{SPRT}_{\text{variance}} \leq \ln(\beta/(1-\alpha))$, then accept hypothesis H_0 as true;
2. If $\text{SPRT}_{\text{variance}} \geq \ln((1-\beta)/\alpha)$, then accept hypothesis H_2 as true; and
3. If $\ln(\beta/(1-\alpha)) < \text{SPRT}_{\text{variance}} < \ln((1-\beta)/\alpha)$, then make no decision and continue sampling.

Each snapshot of residuals (one residual "signal" per component signal from the complex signal) that is passed to the SPRT test module, can have SPRT test decisions for positive mean, negative mean, and variance for each parameter in the snapshot. In an empirical model-based monitoring system according to the present invention, any such SPRT test on any such parameter that results in a hypothesis other than H_0 being accepted as true is effectively an alert on that parameter. Of course, it lies

within the scope of the invention for logic to be inserted between the SPRT tests and the output alerts, such that a combination of a non- H_0 result is required for both the mean and variance SPRT tests in order for the alert to be generated for the parameter, or some other such rule.

5 The output of the SPRT test module 112 will represent a decision for each decomposed signal input, as to whether the estimate is different or the same. These decisions, in turn, can be used to diagnose the state of the process or equipment being monitored. The occurrence of some difference decisions in conjunction with other sameness decisions can be used as an indicator of likely future machine health
10 or process states. The SPRT decisions can be used to index into a diagnostic lookup database, automatically diagnosing the condition of the process or equipment being monitored.

 Generally, any statistical hypothesis test as known by those skilled in the statistical arts can be substituted for the above-described application of SPRT. In
15 addition, decisioning methods known in the art such as fuzzy logic sets and neural networks can be used to render a decision with regard to the sameness or difference of the estimates and the actual values.

 In a setup and implementation phase for the embodiment of the invention shown in FIG. 2, nominal data from the complex signal pick-up is collected prior to
20 monitoring the system. That nominal data is collected under expected operating conditions for the process or equipment and at known, acceptable or desired operating states. The nominal data is decomposed using the same decomposition method that will be used during monitoring, to provide periodically sampled signal component values as signature vectors in the reference library. Each vector
25 represents an observation or snapshot, that is, a time-correlated set of values, one value from each of the decomposed components. Thus, for example, if a complex waveform such as an electrical current signal from a monitored electrical motor is analyzed using the invention, it may be decomposed using frequency filters to provide three frequency-range components, say <58Hz, 58Hz to 62HZ, and >62Hz,

for each of two polarities, providing six inputs to the model. Then the reference library vectors comprise six-element vectors, each element being a signal amplitude at a time-correlated moment from each of the six frequency band component signals.

The collected nominal data is not used in its entirety in the reference library.

5 Rather, the data is distilled down to a representative set sufficient to characterize the population of such snapshots. One method for doing this is described in the aforementioned U.S. Patent No. 5,764,509, and is described therein as Min-Max. Essentially, vectors are selected for inclusion in the reference library if they contain a minimum or a maximum value for any one of the vector elements, across the entire
10 nominal data set. This results in a set of vectors enumerated to no more than twice the number of vector elements. This group of vectors can be augmented using any number of methods with further vectors from the nominal data set.

Figure 3 shows another application of the present invention wherein the monitoring system 120 decodes encoded data from a noisy communications signal,
15 e.g., a QAM or DMT signal. Elements identical to the system of Figure 1 are labeled identically. For such a communications signal, the originally transmitted signal comprises typically a set of superimposed sinusoidal waves, the amplitude and period of which determines the data bits encoded in the carrier. Typically, a specific data word length is encoded together in these superimposed waves, such that
20 decoding the waves at the receiving end results in a data word of the same bit length, e.g., 16 bits. The total "vocabulary" of superimposed waves that the carrier can take on is thus known ahead of time, and constitutes a finite set of states or datagrams (e.g., 16 for a simple QAM). Any deviation from this is noise or degradation present in the transmission line. The embodiment of FIG. 3
25 advantageously provides a computationally efficient way of finding the datagram in the reference library that is most similar to the transmitted signal, and thus most likely the originally intended datagram. So, in this embodiment, instead of directly decoding the received noisy communications signal to extract encoded data, as is done in the art, the noisy signal is compared in the empirical modeling engine 106

with signal signature vectors from reference library 104. The signature vector that is most closely matched by the noisy input signal is identified and the matching signal signature set vector is decoded based on that comparison. The matched vector may be decoded using a look up table 122, for example, and output as the decoded digital
5 result 124.

The output 124 can be the literal data values from the reference library or an index into the signal signature set; or as shown in this example, the output 124 can be an index selecting a decoded value at a location in the lookup table 122. In this example, the similarity operation is carried out in the empirical modeling module
10 106 which compares input snapshots decomposed from the noisy, attenuated or degraded complex input signal with signal signature set snapshots, i.e., signature vectors. For each input snapshot, this comparison identifies the one reference library snapshot that has the greatest measure of similarity to the current state of the input signal, i.e., to the current input snapshot, using the similarity operator
15 described above. A key value is associated with each of the signature vectors and is linked to a lookup table, thus identifying a corresponding decoded digital value with each signature vector. The digital value corresponding to the original signal is output from the lookup table 122 for each match.

To implement the system to be used to extract data from encoded information
20 transmission, e.g. QAM or DMT, the reference library does not need to be populated from empirical data, as is the case for process or machine monitoring. Instead, because the original finite set of datagrams is known ahead of time, the reference library can be constructed from these directly.

Figure 4 shows 16 expected signal states (e.g., encoded transmission signals)
25 conveying information using QAM, for example, each state corresponding to a different one of a four bit digital code. These 16 expected states or signatures are sampled and decomposed into component signals during a signature characterization phase using wavelet analysis, for example. Each state is decomposed into wavelet coefficients that form what are referred to as Detail Levels.

Detail level 1 includes 64 wavelet coefficients; Detail level 2 includes 32; Detail Level 3 includes 16; Detail Level 4 includes 8; and Detail Level 5 includes 4. The wavelet coefficients from Detail Levels 3, 4, and 5 and the Approximation Level 5 yield 32 wavelet coefficients (16, 8, 4 and 4, respectively) and provide sufficient information for characterizing each of the 16 signatures. Thus, the reference library for decoding these datagrams can comprise 16 vectors, each having 32 elements, corresponding to the values of the 32 wavelet coefficients. In operation, the transmission signal (the complex signal being analyzed) is windowed and an appropriate window representing a datagram is decomposed using the same discrete wavelet transform into detail level 3, 4 and 5 and the approximation level 5 to provide a vector of 32 elements. This input vector is compared using the similarity operator in the empirical modeling engine to each of the 16 reference library vectors, and the highest similarity is designated the intended datagram.

FIG. 5 shows an example of discrete wavelet transform of the Signal Number 4 of FIG. 4. In this example, detail levels 1 and 2 are shown, but are not included in the signal signature set or similarity analysis, so that a vector formed to represent the Signal Number 4 in the reference library will contain 32 elements shown in FIG. 5: 16 from the detail level 3, 8 from detail level 4, four from detail level 4, and four from approximation level 5. However, it is understood that the unused wavelet coefficients from detail levels 1 and 2 could be included if desired. Methods are known in the art for developing a full wavelet packet decomposition tree for a complex signal, and ascertaining which detail levels and which approximations contain significant information. These prior art methods can be used in conjunction with the current invention to determine the composition of the vectors that make up the reference library and the input snapshot. The empirical modeling engine is able to render a similarity score for each comparison regardless of the wavelet decomposition used, as long as just one wavelet analysis is selected for processing the input and the reference library.

An example of a noisy or distorted version of each signature from FIG. 4 is shown in FIG. 6. Random additive noise has been added to the signatures in FIG. 6, however the signals have not been attenuated. The decomposition module 102 receives these distorted complex signals - each encoded with one of the 16 signatures - and decomposes them. The decomposition module 102 may be disposed to first identify a window on the complex signal for decomposition, particularly when discrete datagrams are being carried on an essentially continuous carrier wave. A variety of methods are known in the telecommunications arts for achieving this windowing, such as including a start pulse and stop pulse of particular frequency and/or amplitude to identify the window. As an alternative, the decomposition module 102 can continuously decompose the carrier signal over a shifting window of fixed width, and the decomposed components of each such shifted window can be processed for similarity by the empirical modeling engine 106, such that only similarities above a minimum threshold result in identification of a datagram. Other methods known in the art may be used in the present invention for windowing.

The real-time noisy complex signal is periodically decomposed (according to the windowing method chosen) into detail level components to produce a current snapshot of 32 wavelet coefficients, again dispensing with detail levels 1 and 2. The empirical modeling engine 106 compares the detail level components against stored signal signature vectors. So, the empirical modeling engine 106 accepts each snapshot, i.e., the wavelet coefficients in detail levels 3, 4, 5 and approximation level 5, and performs a similarity operation for that snapshot versus each of the 16 signature vectors in the reference library signal signature set. More specifically, in each signature comparison, a pair of sets of 32 wavelet coefficients (one signal signature vector and the decomposed snapshot from the actual input) are compared, comparing each corresponding element in the pair. The comparison identifies the one signature vector of the 16 in the signature set that most closely matches the input snapshot. Then, the empirical modeling engine 106 selects the four-bit code

corresponding to the matched signature vector from the lookup table 122 and outputs that code as a digital output 124.

So, for example, using the bounded angle ratio test (BART) for the comparison, a minimum and a maximum value at each end of a right triangle
5 hypotenuse defines the range found for a given coefficient across all 16 signatures in the reference library. The input signal coefficients are compared for similarity, mapping each input coefficient and a corresponding signature vector coefficient along the particular hypotenuse according to the value of the coefficient within the range. Line segments are drawn from each mapped point to the apex. The internal
10 angle between the drawn line segments is less than or equal to ninety degrees. That internal angle is normalized to 90 degrees and, the normalized result is subtracted from one to yield a similarity value between one and zero. That similarity value is a measure of similarity for the pair of corresponding elements, one (1) indicating identity and zero (0) indicating dissimilarity. The similarity values for all 32
15 components (i.e. wavelet coefficients) in each pair are averaged to provide an overall measure of similarity between the signatures.

Turning to FIG. 7, results are shown of similarity measurements between the noisy signals of FIG. 6 and the original datagrams of FIG. 4. Each of the 16 graphs corresponds to one of the noise-distorted signal graphs in FIG. 6. Within each graph
20 are sixteen bars comprising the bar chart, one bar for each of the 16 model signals from FIG. 4 that make up the reference library. Each bar represents the vector similarity value of a reference library datagram to the particular noisy datagram from FIG. 6 for that chart. In each comparison, one reference library signature is identified as being closest to the noisy input vector. The match is indicated by an
25 asterisk on the bar having the highest similarity value. In each of the 16 comparisons, the preferred embodiment system automatically scored the correct noisy input signal as being the most similar to the original reference library signature.

Accordingly, the system of the present invention can accurately identify the correct reference signature from a noisy input signature, thereby, acting as a superior filter for transmission noise to facilitate extracting correct information content from a degraded carrier signal. The reference library may also have associated with each vector an index for directing a lookup into a lookup table, so that encoded digital information may be automatically extracted and decoded with a much higher degree of accuracy than has been heretofore achieved.

FIG. 8 is a flow diagram 130 showing the steps of the preferred embodiment of the present invention. First in step 132, a signal is received and passed to the decomposer. Then, in step 134, the signal is decomposed into components using, for example, wavelet decomposition or signal filters. Next, in step 136, the decomposed signal components are compared against the signal signature set. In step 138, a check is made to determine if additional steps are to be taken or if additional analysis is to be done on the comparison results and, if not, in step 140 the comparison result is output.

If additional steps are to be taken on the comparison result, then, in step 142 a check is made to determine whether data is to be extracted from the signal, e.g., communication data. If it is determined that the signal contains encoded data (e.g. the signal is a data transmission signal), then, in step 144, the signal signature vector closest to the input snapshot is identified. In step 146 the lookup table is searched for a digital value corresponding to the matching vector and, that corresponding digital value is passed to the output.

If the signal is not carrying encoded data but, instead is a signal from a sensor monitoring a system, then, in step 150, a G matrix is determined, being equivalent to the first term from equation 7 above:

$$\underline{\hat{W}} = \overline{G}^{-1} \bullet \left(\overline{D}^T \otimes \vec{y}_{input} \right) \quad (16)$$

$$\overline{G} = \left(\overline{D}^T \otimes \overline{D} \right) \quad (17)$$

It should be noted that the G matrix may be generated prior to receiving any signal, then stored locally, retrieved for use in step 150 and made available for subsequent system analysis. In step 152 the expected signal is generated by multiplying the G matrix with the similarity values vector, as per equation 5 above. In step 154 the difference between the expected signal and the input signal is calculated. In step 156 SPRT analysis is applied to the calculated result from step 154. Finally, in step 158, an appropriate diagnostic is run on the SPRT result.

It is contemplated that the present invention has wide application in a number of fields, and may take many physical embodiments. In any process control setting, the invention can embody a computer connected to a network or the process control system, wherein the complex signal would be received by means of a data acquisition card or a network card, and processed in the computer processor according to software instructions in memory. The output of the computer embodiment can be supplied to the screen of the computer, or made available via a network connection to either downstream processing, such as in the control system itself, or to remote viewing, as in a web browser on another computer. In another embodiment, the invention comprises a hardware device specifically for analyzing the complex signal, which may be a heartbeat or an electric motor current. In this embodiment, a processor is also typically provided in the device, along with memory for storing the data and the model or reference data. The device then preferably includes the sensor or set of sensors and their electrical signal connections for detecting the complex signal in question. In an on-board embodiment, the invention may comprise a set of microprocessor instructions that are part of a broader digital signal processor (DSP) instruction set running a DSP embedded in a machine, such as an automobile or an electric motor. In this case, receipt of the complex signal is typically already provided for in the embedded DSP as an existing sensor, such as a voltmeter or the like. Finally, the invention can also be practiced off-line as a program in a standard desktop or workstation environment, where the complex signal is provided as a data file stored on media. Generally across these

embodiments, the signal decomposition module can be a software-driven processor, a DSP or other microprocessor dedicated to the task, an application-specific integrated circuit (ASIC) customized to decompose in hardware, or even an arrangement of simple circuit elements (such as filters) that decompose the complex
5 signal in its analog state, prior to any digitization for processing.

It should be appreciated that a wide range of changes and modifications may be made to the embodiments of the invention as described herein. Thus, it is intended that the foregoing detailed description be regarded as illustrative rather than limiting and that the following claims, including all equivalents, are intended
10 to define the scope of the invention.